

How to Write a Philosophy Paper

The goal of a philosophy paper is simple: make a compelling argument. This guide aims to teach you how to write philosophy papers, starting from the ground up. To do that, we need to understand *what an argument is* and *what makes an argument compelling*.

1. What's an argument?

An **argument** is a combination of **premises** and a **conclusion**. Here's an example:

- (a) Harper drives a Maserati, so Harper is rich.

Here *Harper drives a Maserati* serves as the premise for the conclusion *Harper is rich*. Another way we can write this argument is in premise-conclusion form:

- (b) Premise #1: Harper drives a Maserati.
Conclusion: Thus, Harper is rich.

And if we found this argument in everyday speech, it might look like this:

- (c) Woah, did you see who was driving that Maserati? Harper! She must be rich.

Arguments, then, can appear in many forms, but they always consist of at least one premise and a conclusion. Think of premises as *reasons* that aim to show the conclusion is true. In our example, the reason we ought to believe Harper is rich (the argument alleges) is that she drives a Maserati.

2. What makes an argument compelling?

Roughly, a compelling argument is an argument that gives convincing reasons to believe its conclusion. Let's consider our argument. Should the fact that Harper drives a Maserati convince us that she is rich?

Let's think about it. Maserati's are expensive cars, no doubt. But for all we know—that's to say, for all the argument tells us—Harper may not even own the Maserati; perhaps she is borrowing it from her rich friend Makayla. Or perhaps Harper *does* own the Maserati, but purchasing it caused her to go bankrupt. Since both cases are possible,

it does not seem like the fact that Harper drives a Maserati is a convincing reason to believe she is rich.

Note #1: Philosophers would call the fact that Harper drives a Maserati a **prima facie reason** to believe she is rich. Prima facie reasons provide *some* reason to believe a conclusion, but they do not provide *sufficient* reason to believe that conclusion.

We can evaluate arguments more precisely using the notions of **validity** and **soundness**.

2a. Validity

An argument is **valid** if its conclusion *must* be true if its premises are true, and an argument is **invalid** if it is not valid.

That may sound complicated, but actually testing for validity is straightforward: assume all the premises are true, and ask whether the conclusion can be false. If the answer is yes, the argument is invalid. If the answer is no, the argument is valid.

Let's try it out with a new argument.

- (d) Premise #1: Alyssa thinks euthanasia is morally permissible.
Premise #2: Taylor thinks euthanasia is morally impermissible.
Conclusion: Thus, there is no right answer as to whether euthanasia is morally permissible or impermissible.

Is the argument valid?

Assume that both premises are true. Can the conclusion still be false?

Yes. *Even if* Alyssa and Taylor disagree about euthanasia, it can still be the case that there is a right answer as to whether euthanasia is morally permissible. Thus, since the conclusion can be false even when both premises are true, the argument is invalid.

Note #2: Validity exclusively concerns the *form* of the argument. So, we can test for validity by retaining the same *form* of argument but substituting for its *content*. This is often beneficial, because it often makes testing for validity easier. Let's try it out with our argument.

- (e) Premise #1: Alyssa thinks x is y .
Premise #2: Taylor thinks x is not y .

Conclusion: Thus, there is no right answer as to whether x is y or not y .

See how that preserves the same form of argument? All we've done is substitute variables for 'euthanasia,' 'morally permissible,' and 'morally impermissible.' Now we can assign different content to those variables to see whether this form of argument is valid. For example:

- (f) Premise #1: Alyssa thinks the earth is flat.
Premise #2: Taylor thinks the earth is not flat.
Conclusion: Thus, there is no right answer as to whether the earth is flat or not flat.

Even if Alyssa and Taylor disagree about whether the earth is flat, there is still a right answer to the question! This form of argument is invalid.

Note #3: An argument can be valid even if its premises are false, its conclusion is false, or both its premises and conclusion are false. Consider the following argument:

- (g) Premise #1: If lemons are sweet, then robots have souls.
Premise #2: Lemons are sweet.
Premise #3: So, robots have souls.

Obviously, neither premise nor the conclusion of this argument is true. But is the argument valid? Remember, when testing for validity simply *assume* the premises are true, and ask whether in that case the conclusion can possibly be false. If the answer is no, the argument is valid. This argument is valid!

2b. Soundness

Besides validity, the other way to precisely evaluate an argument is to consider whether it is sound. An argument is **sound** if it is valid and all its premises are true, and an argument is **not sound** if it is invalid or if one of its premises is false. When testing for whether an argument is sound, then, first test for whether it is valid, and then consider whether all of its premises are true.

Let's try it out with a canonical argument example.

- (h) Premise #1: Socrates is a human.
Premise #2: All humans are mortal.
Conclusion: Therefore, Socrates is mortal.

Is the argument sound? Notice that it is impossible for its premises to be true but its conclusion to be false (so, the argument is valid). Notice, too, that both of its premises

are true. So, since the argument is valid and all of its premises are true, the argument is sound!

3. How do I extract an argument from a text?

It is easiest to evaluate arguments when they are in premise-conclusion form, but rarely do philosophers present their arguments in premise-conclusion form in their published writings. One helpful skill to have, then, is the ability to extract an argument from a text and analyze it into its premises and conclusion. After you have done so, you can easily consider whether or not the argument is valid and sound, and thus whether or not it is a compelling argument.

Let's try it out with a slightly revised argument from Descartes' *Meditations*.

Descartes: We should not think that non-human animals speak. For if that were true, then they could make themselves understood by us, but we cannot understand them.

Can you identify the premises and conclusion of Descartes' argument?

Here are two more arguments for practice. The first is adapted from Epicurus' *Letter to Menoeceus*, and the second is a Kanye West quotation.

Epicurus: Everyone would agree that you should not fear something that is not at all bad, so get used to believing that death is nothing to be feared. Only pain is bad, and when you are dead you cannot feel pain.

Kanye: Let me give my definition of crazy. Webster's tells that the definition of crazy is to do the same thing expecting a different result. My definition of crazy is to think of something that nobody would ever do and attempt to do it. And while people are looking at it and it's not working people will say, 'Oh, that's crazy.' But if you do something that no one would ever do and it works then that's genius. So I ought to be crazy because crazy is the first step to genius.

Tips for extracting an argument:

First, identify the *conclusion*.

- Conclusions most often come in the form of *assertions*, and they are often indicated by a word like ‘thus,’ ‘so,’ ‘therefore,’ or ‘then.’

Next, identify the reasons the argument gives to support the conclusion. These will be the argument’s *premises*.

- It can be helpful at first to simply copy down the exact words the author uses. This will give you a sense of how the argument works, and then you can polish it up afterward to make it more streamlined.
- Do not expect *everything* an author says to be a premise! Also, it is best to analyze an argument into *as few premises as possible*.
- Occasionally, some premises will be left *implied* or *assumed*. So, do not think that *only* premises that are explicitly stated are used for the argument!

4. Okay, but where do I start when writing my essay?

Most often, and especially in philosophy courses at the introductory level, essay prompts will ask you to write some variation of a common theme:

Is Philosopher So-and-so right in claiming that such-and-such? In giving your answer, be sure to respond to So-and-so’s argument for thinking that such-and-such is true.

There is a very straightforward method for starting to outline such an essay—and now that we’ve gone through the previous sections, you will be able to carry it out.

The first step is to reconstruct the author’s argument for the claim, analyzing it into premise-conclusion form (see section 3). The next step is to evaluate the argument (see section 2). Is the argument valid and sound? If not, then you should dedicate part of your paper to showing that the argument fails.

Let's imagine an essay prompt of this form that concerns Epicurus' argument above:

Is Epicurus right in claiming that death is nothing to be feared? In giving your answer, be sure to respond to Epicurus' argument.

Our first task in planning the essay is to reconstruct Epicurus' argument:

Premise #1: You should not fear something that is not at all bad for you.

Premise #2: Only pain is bad for you.

Premise #3: When you are dead, you cannot feel pain.

Conclusion: Thus, you should not fear death.

With the argument in this form, we can easily think about whether the argument is compelling. Notice that the argument seems valid: if the premises are true, it does seem that the conclusion must be true. So it doesn't look like we'll be able to argue that the reasoning behind the argument is faulty. But is the argument sound? Are all the premises true?

Let's assume premise #3 is true. What about premises #1 and #2?

Well, it's tricky, right? Think about premise #2: *Only pain is bad*. A good way to proceed is to ask whether there is anything that is bad but not painful (philosophers would call that a **counterexample** to the claim that *only pain is bad*). Maybe loss of opportunity is bad for you, even if it is not painful? If so, then premise #2 would not be true, and the argument would fail. You could then write your essay arguing that Epicurus is not right to claim that you should not fear death, because he is wrong to think that only pain is bad for you: loss of opportunity is also bad for you, even though it is not painful, and since death takes away opportunity, death is bad for you.

5. Okay, but how do I structure my essay?

Structuring a philosophy essay is simple, so long as you know exactly what you are going to argue.

Introductory paragraph

Your introduction has one task: to state what you are going to argue and briefly indicate how you will argue for it. Often this task is aided by a preliminary sentence or two that sums up the position of your main interlocutor(s) in the essay. For example:

Epicurus argues that we have no reason to fear death, in part because he thinks that only pain is bad. I argue against Epicurus that loss of opportunity is bad, even if it is not painful, and thus that we do have reason to fear death.

The best introductions, like this one, are straightforward and concise. No word is wasted, and no unnecessary information is included.

First body paragraph

The task of your first paragraph is to provide the necessary background for your argument. Often this will consist of summing up the positions of your interlocutor(s) in further detail. The main thing to remember here is that you should not include any information that is not *absolutely necessary* for the argument you will go on to make. Here's an example of how a first paragraph might look for our essay on Epicurus:

In *Letter to Menoecus*, Epicurus argues that we should not fear death. His main thought is a reasonable one: we should not fear something that is not bad for us. If, then, death is something that is not bad for us, Epicurus would be right in claiming that we should not fear death. So, is death bad for us, or is it not? Epicurus thinks it is not. For Epicurus, only pain is bad, and when we are dead we cannot feel pain. If both those claims are true, then death would not be bad for us, and so we should not fear it. Let's assume that Epicurus is right in thinking that when we are dead we cannot feel pain (some may disagree, but it would be difficult to prove the matter one way or the other). Is Epicurus right in claiming that only pain is bad for us?

No, I will argue. For consider...

The best first paragraphs, like this one, do three things. They *explain* how your main interlocutor's argument works. They *clarify* exactly what you think is reasonable and what you want to call into question. And they *identify* the key question that will set the agenda for the rest of the essay.

Second body paragraph onward

You should strive to begin your own argument in the second paragraph. If we were to proceed with this essay, the second paragraph is where we would begin arguing for the claim that loss of opportunity, even though it is not painful, can also be bad for us. We might use one example to help make our point. We might also entertain a strong objection to that claim, and then give an argument that defeats the objection. As the essay goes on, we would also want to argue that death causes a loss of opportunity for us, and thus causes something bad for us, and thus, contrary to what Epicurus claims, we do have reason to fear death. In short, the bulk of your essay should be devoted to giving your own argument, and the best place to begin that is in the second paragraph after your introduction.

Conclusion

Most students are tempted to repeat the main lines of their argument in their conclusion. This is not a bad strategy (and may be welcome if your argument involves many steps). Alternatively, you can use your conclusion to emphasize the coolest thing that came up in your essay, or to muse a bit on why your main point matters.

6. Tips

(i) Focus your paper on *one* small but important point

The best philosophy papers, especially at the undergrad level, identify the *one* point on which some dispute turns and then focus their argument (entirely) on that one point. Strive to do that. As much as you can, avoid the temptation to make your argument about some big, far-reaching claim.

(ii) Develop *one* argument for your main claim

It is far better to give one strong and fully developed argument than several underdeveloped arguments. After all, if the one argument proves your point, why do we need to hear the other arguments? An exception to this rule is if there is an obvious argument that supports your main claim, but actually it is problematic, and this other argument is better. In that case, you would want to *briefly* give the

obvious argument, then *briefly* show that it is problematic, then devote the bulk of your paper to the other argument.

(iii) Limit quotations

As a general rule, you should use a quotation only when an author's *phrasing* is extremely important. If you are simply summing up an author's position, it is much better to do that in your own words than with a quotation. Also, if you do give a quotation, *be sure to explain what you think it means* in the sentence immediately after the quotation. Otherwise, your reader cannot be confident that she and you interpret the quotation in the same way.

7. Five most common writing mistakes

(i) Including unnecessary information

Often students think they need to explain everything an author says about some topic. This is a mistake. You should include only information that is necessary to set up your own argument, or that shows the reader how your argument relates to what others have said.

(ii) Not using signposts

Every sentence has a purpose, e.g. to make a claim, to make an objection, to give a reason, to explain, to give an example. Challenge yourself to identify the purpose behind each of your sentences, and be sure to make clear to the reader what you intend for the sentence to do. The best way to do this is to use **signposts** in your writing. A list of common signposts: *thus, so, therefore, then, the upshot is* (to indicate a conclusion); *because, the reason is, since* (to give a reason); *for example, for instance* (to indicate an example); *further, additionally, moreover* (to continue a thought); *but, however* (to introduce an objection).

(iii) Clunky writing

Strive to write clearly and concisely. Keep your sentences short. Do not use a sophisticated vocabulary. Choose the simplest words to express what you want to say. Limit dependent clauses. Note: it is actually *harder* to write this way than to write in the opposite way, but it will make your writing stronger.

(iv) Irresponsible pronoun use

Try to limit your use of pronouns, especially ‘it,’ ‘this,’ and ‘that.’ The reason is that, unless pronouns are used very skillfully, it is often unclear to what exactly they are intended to refer. Forcing yourself to avoid these words will clarify in your own mind exactly what it is you are trying to say.

(v) Not giving reasons for your claims

Anytime you make a claim that is controversial, you should give a reason for why you think that claim is true. Always keep in mind that philosophy is a game of giving reasons: the claim with the best reasons behind it wins.

8. Bad writing example

Often the best way to learn is by example. Here, then, is an example of a bad introductory paragraph (written by me to be intentionally bad). Try to find the mistakes with each sentence. Hint: there is only one sentence that would survive if we were to revise the paragraph to be good.

In this passage Aquinas discusses various aspects of cognition. Aquinas says, “It is entirely impossible for all human beings to share in a single intellect.” He then goes on to give three possibilities for how the intellect is joined to a human. But what does Aquinas mean by ‘intellect’? What an intellect *is* is a very controversial topic and many people even today disagree about it (for example, brain vs. mind debate), but Aquinas assumes he knows what it is. Nevertheless there are good points in his argument and it behooves us to inspect his postulations with, as Shakespeare says, ‘hawkish eye and cutting mind.’ The main reason Aquinas thinks it is impossible for humans to share a single intellect is that, if we did, there would be only one action of thinking, and Aquinas thinks that result is inconsistent with one of our most fundamental experiences: experiencing that it is oneself who thinks.

9. Good writing example

Contrast the above bad writing sample with the good writing sample below, taken from University of Toronto professor Phil Clark's paper "Inescapability and the Analysis of Agency":

A strategy is for doing something. What does Velleman want the Kantian strategy to do? The answer, I think, is that he wants it to resolve an apparent tension within his view. He wants to "explain how morality can be objective." And he does this by first explaining how practical reason can be objective, and then explaining how practical reason supports morality. The problem is that he also accepts claims that can seem to rule out any sort of objectivism about practical reason. He takes these claims from Bernard Williams:

<quotation>.

I agree with Williams's premise that reasons for acting must be able to engage a motive that the agent has or could come to have through sound deliberation; and I do not wish to question the assumption that deliberation can convey him only from motive to motive, so that his current motives determine where he could rationally end up. But I reject Williams's conclusion, that reasons must therefore be geared to something subjective in the agent's psychological make-up (119-20). Velleman grants that nothing can be a reason for action unless it either engages a current motive of the agent or engages a motive that could be derived from a current motive of the agent. He thus accepts that reasons must be "geared to" something in the agent's psychological make-up; they must be geared to some current motive of the agent. And this seems to spell doom for any thought that the demands of practical reason, and of any morality spun from them, might be objective. For surely an objective demand would be one that was there in the world quite apart from what the agent happens to want. So here is the problem. If practical reason is tethered to current motives of the agent, how can it be objective?